

Controllable-choice Message Sequence Graphs

Martin Chmelík^{1*} and Vojtěch Řehák^{2**}

¹ Institute of Science and Technology Austria (IST Austria), Klosterneuburg, Austria
martin.chmelik@ist.ac.at

² Faculty of Informatics Masaryk University, Brno, Czech Republic
rehak@fi.muni.cz

Abstract. We focus on the realizability problem of Message Sequence Graphs (MSG), i.e. the problem whether a given MSG specification is correctly distributable among parallel components communicating via messages. This fundamental problem of MSG is known to be undecidable. We introduce a well motivated restricted class of MSG, so called controllable-choice MSG, and show that all its models are realizable and moreover it is decidable whether a given MSG model is a member of this class. In more detail, this class of MSG specifications admits a deadlock-free realization by overloading existing messages with additional bounded control data. We also show that the presented class is the largest known subclass of MSG that allows for deadlock-free realization.

1 Introduction

Message Sequence Chart (MSC) [15] is a popular formalism for specification of distributed systems behaviors (e.g. communication protocols or multi-process systems). Its simplicity and intuitiveness come from the fact that an MSC describes only exchange of messages between system components, while other aspects of the system (e.g. content of the messages and internal computation steps) are abstracted away. The formalism consists of two types of charts: (1) basic Message Sequence Charts (bMSC) that are suitable for designing finite communication patterns and (2) High-level Message Sequence Charts (HMSC) combining bMSC patterns into more complex designs. In this paper, we focus on the further type reformulated as Message Sequence Graphs (MSG) that has the same expressive power as HMSC but a simpler structure, and hence it is often used in theoretical computer science papers, see, e.g. [2,4,6,13,22].

Even such incomplete models as MSG can indicate serious errors in the designed system. The errors can cause problems during implementation or even make it impossible. Concerning verification of MSC models, researchers have studied a presence of a race condition in an MSC [3,7,10,22], boundedness of

* The author was supported by Austrian Science Fund (FWF) Grant No P 23499-N23 on Modern Graph Algorithmic Techniques in Formal Verification, FWF NFN Grant No S11407-N23 (RiSE), ERC Start grant (279307: Graph Games), and Microsoft faculty fellows award.

** The author is supported by the Czech Science Foundation, grant No. P202/12/G061.

the message channels [4], the possibility to reach a non-local branching node [6,19,13,16,17,11,20], deadlocks, livelocks, and many more. For a recent overview of current results see, e.g. [9].

In this paper, we focus on the realizability problem of MSG specifications, i.e. implementation of the specification among parallel machines communicating via messages. This problem has been studied in various settings reflecting parameters of the parallel machines, the environment providing message exchanges as well as the type of equivalence considered between the MSG specification and its implementation. Some authors restricted the communication to synchronous handshake [13,12], needed several initial states in the synthesized machines [5], or considered language equivalence with global accepting states in the implementation (the implementation accepts if the components are in specific combinations of its states) [21]. From our point of view, the crucial aspect is the attitude to non-accepted executions of the implementation. When language equivalence is taken into account, an intentional deadlock can prevent a badly evolving execution from being accepted [13]. In our setting every partial execution can be extended into an accepting one. Therefore, we focus on a deadlock-free implementation of a given MSG into Communicating Finite-State Machines (CFM) with FIFO communicating channels and distributed acceptance condition, i.e. a CFM accepts if each machine is in an accepting state. In [18], it has been shown that existence of a CFM realizing a given MSG without deadlocks is undecidable. When restricted to *bounded* MSG (aka regular MSG, i.e. communicating via finite/bounded channels, and so generating a regular language), the problem is EXPSPACE-complete [18].

In later work [13,5], a finite data extension of messages was considered when realizing MSG. This is a very natural concept because message labels in MSG are understood as message types abstracting away from the full message content. Hence, during implementation, the message content can be refined with additional (finite) data that helps to control the computation of the CFM in order to achieve the communication sequences as specified in the given MSG. The main obstacle when realizing MSG are nodes with multiple outgoing edges — choice nodes. In a CFM realization, it is necessary to ensure that all Finite-State Machines choose the same successor of each choice node. This can be hard to achieve as the system is distributed.

In [13], a class of so called *local-choice* MSG [17,6] was shown to include only MSG realizable in the above mentioned setting. Local-choice specifications have the communication after each choice node initiated by a single process — the choice leader. Intuitively, whenever a local-choice node is reached, the choice leader machine attaches to all its outgoing messages the information about the chosen node. The other machines pass the information on. This construction is sufficient to obtain a deadlock-free realization, for details see [13]. Another possible realization of local-choice MSG is presented in [16]. Due to [11], it is also decidable to determine whether a given MSG is language equivalent to some local-choice MSG and, moreover, each equivalent MSG can be algorithmically realized by a CFM. To the best of our knowledge, this is the largest

class of deadlock-free realizable specifications in the standard setting, i.e. with additional data, FIFO channels, and local accepting states.

In this paper, we introduce a new class of *controllable-choice* MSG that extends this large class of realizable MSG. The crucial idea of controllable-choice MSG is that even some non-local-choice nodes can be implemented, if the processes initiating the communication after the choice can agree on the successor node in advance. This is achieved by exchanging bounded additional content in existing messages. We call choice nodes where such an agreement is possible *controllable-choice* nodes, and show that the class of MSG with these nodes is more expressive than the class of MSG that are language equivalent to local-choice MSG.

2 Preliminaries

In this section, we introduce the Message Sequence Chart (MSC) formalism that was standardized by the International Telecommunications Union (ITU-T) as Recommendation Z.120 [15]. It is used to model interactions among parallel components in a distributed environment. First, we introduce the basic MSC.

basic Message Sequence Charts (bMSC) Intuitively, a bMSC identifies a single finite execution of a message passing system. Processes are denoted as vertical lines — instances. Message exchange is represented by an arrow from the sending process to the receiving process. Every process identifies a sequence of actions — sends and receives — that are to be executed in the order from the top of the diagram. The communication among the instances is not synchronous and can take arbitrarily long time.

Definition 1. A basic Message Sequence Chart (bMSC) M is defined by a tuple $(E, <, \mathcal{P}, \mathcal{T}, P, \mathcal{M}, l)$ where:

- E is a finite set of events,
- $<$ is a partial ordering on E called visual order,
- \mathcal{P} is a finite set of processes,
- $\mathcal{T} : E \rightarrow \{\text{send}, \text{receive}\}$ is a function dividing events into sends and receives,
- $P : E \rightarrow \mathcal{P}$ is a mapping that associates each event with a process,
- $\mathcal{M} : \mathcal{T}^{-1}(\text{send}) \rightarrow \mathcal{T}^{-1}(\text{receive})$ is a bijective mapping, relating every send with a unique receive, such that a process cannot send a message to itself, we refer to a pair of events $(e, \mathcal{M}(e))$ as a message, and
- l is a function associating with every message (e, f) a label m from a finite set of message labels \mathcal{C} , i.e. $l(e, f) = m$.

Visual order $<$ is defined as the reflexive and transitive closure of $\mathcal{M} \cup \bigcup_{p \in \mathcal{P}} <_p$ where $<_p$ is a total order on $P^{-1}(p)$.

We require the bMSC to be *first-in-first-out (FIFO)*, i.e., the visual order satisfies for all messages $(e, f), (e', f')$ and processes p, p' the following condition

$$e <_p e' \wedge P(f) = P(f') = p' \Rightarrow f <_{p'} f'.$$

Every event of a bMSC can be represented by a letter from an alphabet

$$\Sigma = \{p!q(m) \mid p, q \in \mathcal{P}, m \in \mathcal{C}\} \cup \{q?p(m) \mid p, q \in \mathcal{P}, m \in \mathcal{C}\}.$$

Intuitively, $p!q(m)$ denotes a send event of a message with a label m from a process p to a process q , and $q?p(m)$ represents a receive event of a message with a label m by q from a process p . We define a *linearization* as a word over Σ representing a total order of events that is consistent with the partial order $<$. For a given bMSC M , a *language* $\mathcal{L}(M)$ is the set of all linearizations of M .

Message Sequence Graphs It turns out that specifying finite communication patterns is not sufficient for modelling complex systems. Message Sequence Graphs allow us to combine bMSCs into more complex systems using alternation and iteration. An MSG is a directed graph with nodes labeled by bMSCs and two special nodes, the initial and the terminal node. Applying the concept of finite automata [14], the graph represents a set of paths from the initial node to the terminal node. In MSG, every such a path identifies a sequence of bMSCs. As every finite sequence of bMSCs can be composed into a single bMSC, an MSG serves as a finite representation of an (infinite) set of bMSCs.

Definition 2. A Message Sequence Graph (MSG) is defined by a tuple $G = (\mathcal{S}, \tau, s_0, s_f, L)$, where \mathcal{S} is a finite set of states, $\tau \subseteq \mathcal{S} \times \mathcal{S}$ is an edge relation, $s_0 \in \mathcal{S}$ is the initial state, $s_f \in \mathcal{S}$ is the terminal state, and $L : \mathcal{S} \rightarrow \text{bMSC}$ is a labeling function.

W.l.o.g., we assume that there is no incoming edge to s_0 and no outgoing edge from s_f . Moreover, we assume that there are no nodes unreachable from the initial node and the terminal node is reachable from every node in the graph.

Given an MSG $G = (\mathcal{S}, \tau, s_0, s_f, L)$, a *path* is a finite sequence of states $s_1 s_2 \dots s_k$, where $\forall 1 \leq i < k : (s_i, s_{i+1}) \in \tau$. A *run* is defined as a path with $s_1 = s_0$ and $s_k = s_f$.

Intuitively, two bMSCs can be composed to a single bMSC by appending events of every process from the latter bMSC at the end of the process from the precedent bMSC. Formally, the *sequential composition* of two bMSCs $M_1 = (E_1, <_1, \mathcal{P}, \mathcal{T}_1, P_1, \mathcal{M}_1, l_1)$ and $M_2 = (E_2, <_2, \mathcal{P}, \mathcal{T}_2, P_2, \mathcal{M}_2, l_2)$ such that the sets E_1 and E_2 are disjoint (we can always rename events so that the sets become disjoint), is the bMSC $M_1 \cdot M_2 = (E_1 \cup E_2, <, \mathcal{P}, \mathcal{T}_1 \cup \mathcal{T}_2, P_1 \cup P_2, \mathcal{M}_1 \cup \mathcal{M}_2, l_1 \cup l_2)$, where $<$ is a transitive closure of $<_1 \cup <_2 \cup \bigcup_{p \in \mathcal{P}} (P_1^{-1}(p) \times P_2^{-1}(p))$. Note that we consider the weak concatenation, i.e. the events from the latter bMSC may be executed even before some events from the precedent bMSC.

Now, we extend the MSG labeling function L to paths. Let $\sigma = s_1 s_2 \dots s_n$ be a path in MSG G , then $L(\sigma) = L(s_1) \cdot L(s_2) \cdot \dots \cdot L(s_n)$. For a given MSG G , the *language* $\mathcal{L}(G)$ is defined as $\bigcup_{\sigma \text{ is a run in } G} \mathcal{L}(L(\sigma))$. Hence, two MSG are said to be *language-equivalent* if and only if they have the same languages.

Communicating Finite-State Machines A natural formalism for implementing bMSCs are Communicating Finite-State Machines (CFM) that are used for example in [3,1,13]. The CFM consists of a finite number of finite-state machines that communicate with each other by passing messages via unbounded FIFO channels.

Definition 3. *Given a finite set \mathcal{P} of processes and a finite set of message labels \mathcal{C} , the Communicating Finite-State Machine (CFM) \mathcal{A} consists of Finite-State Machines (FSMs) $(\mathcal{A}_p)_{p \in \mathcal{P}}$. Every \mathcal{A}_p is a tuple $(\mathcal{S}_p, A_p, \rightarrow_p, s_p, F_p)$, where:*

- \mathcal{S}_p is a finite set of states,
- $A_p \subseteq \{p!q(m) \mid q \in \mathcal{P}, m \in \mathcal{C}\} \cup \{p?q(m) \mid q \in \mathcal{P}, m \in \mathcal{C}\}$ is a set of actions,
- $\rightarrow_p \subseteq \mathcal{S}_p \times A_p \times \mathcal{S}_p$ is a transition relation,
- $s_p \in \mathcal{S}_p$ is the initial state, and
- $F_p \subseteq \mathcal{S}_p$ is a set of local accepting states.

We associate an unbounded FIFO error-free channel $\mathcal{B}_{p,q}$ with each pair of FSMs $\mathcal{A}_p, \mathcal{A}_q$. In every configuration, the content of the channel is a finite word over the label alphabet \mathcal{C} .

Whenever an FSM \mathcal{A}_p wants to send a message with a label $m \in \mathcal{C}$ to \mathcal{A}_q , it enqueues the label m into channel $\mathcal{B}_{p,q}$. We denote this action by $p!q(m)$. Provided there is a message with a label m in the head of channel $\mathcal{B}_{p,q}$, the FSM \mathcal{A}_q can receive and dequeue the message with the label m . This action is represented by $q?p(m)$. A *configuration* of a CFM $\mathcal{A} = (\mathcal{A}_p)_{p \in \mathcal{P}}$ is a tuple $C = (s, B)$, where $s \in \prod_{p \in \mathcal{P}} (\mathcal{S}_p)$ and $B \in (\mathcal{C}^*)^{\mathcal{P} \times \mathcal{P}}$ — local states of the FSMs together with the contents of the channels. Whenever there is a configuration transition $C_i \xrightarrow{a_i} C_{i+1}$, there exists a process $p \in \mathcal{P}$ such that the FSM \mathcal{A}_p changes its local state by executing action $a_i \in A_p$ and modifies the content of one of the channels.

The CFM execution starts in an *initial configuration* $s_0 = \prod_{p \in \mathcal{P}} \{s_p\}$ with all the channels empty. The CFM is in an *accepting configuration*, if every FSM is in some of its final states and all the channels are empty. We will say that a configuration is a *deadlock*, if no accepting configuration is reachable from it. A CFM is *deadlock-free* if no deadlock configuration is reachable from the initial configuration. An *accepting execution* of a CFM \mathcal{A} is a finite sequence of configurations $C_1 \xrightarrow{a_1} C_2 \xrightarrow{a_2} \dots \xrightarrow{a_{n-1}} C_n$ such that C_1 is the initial configuration and C_n is an accepting configuration. The word $a_1 a_2 \dots a_{n-1}$ is then an *accepted word* of \mathcal{A} . Given a CFM \mathcal{A} , the *language* $\mathcal{L}(\mathcal{A})$ is defined as the set of all accepted words of \mathcal{A} .

3 Controllable-choice Message Sequence Graphs

For a given MSG we try to construct a CFM such that every execution specified in the MSG specification can be executed by the CFM and the CFM does not introduce any additional unspecified execution.

Definition 4 ([1]). *An MSG G is realizable iff there exists a deadlock-free CFM \mathcal{A} such that $\mathcal{L}(G) = \mathcal{L}(\mathcal{A})$.*

One of the most natural realizations are projections. A projection of a bMSC M on a process p , denoted by $M|_p$, is the sequence of events that are to be executed by the process p in M . For every process $p \in \mathcal{P}$, we construct a FSM \mathcal{A}_p that accepts a single word $M|_p$. This construction is surprisingly powerful and models all of the bMSC linearizations.

Proposition 1. *Let M be a bMSC, then CFM $\mathcal{A} = (M|_p)_{p \in \mathcal{P}}$ is a realization, i.e. $\mathcal{L}(M) = \mathcal{L}(\mathcal{A})$.*

It turns out that the main obstacle when realizing MSG are nodes with multiple outgoing edges — choice nodes. It is necessary to ensure that all FSMs choose the same run through the MSG graph. This can be hard to achieve as the system is distributed.

In what follows, we present a known class of local-choice MSG specifications that admits a deadlock-free realization by adding control data into the messages. Then, we define a new class of *controllable-choice* MSG and compare the expressive power of the presented classes.

Local-choice MSG is a class studied by many authors [6,19,13,16,17,11]. Let M be a bMSC, we say that a process $p \in \mathcal{P}$ *initiates* the bMSC M if there exists an event e in M , such that $P(e) = p$ and there is no other event e' in bMSC M such that $e' < e$. For a given MSG, every node $s \in \mathcal{S}$ identifies a set $\text{triggers}(s)$, the set of processes initiating the communication after the node s . Note that it may not be sufficient to check only the direct successor nodes in the MSG.

Definition 5. *Let $G = (\mathcal{S}, \tau, s_0, s_f, L)$ be an MSG. For a node $s \in \mathcal{S}$, the set $\text{triggers}(s)$ contains process p if and only if there exists a path $\sigma = \sigma_1 \sigma_2 \dots \sigma_n$ in G such that $(s, \sigma_1) \in \tau$ and p initiates bMSC $L(\sigma)$.*

Definition 6. *A choice node u is a local-choice node iff $\text{triggers}(u)$ is a singleton. An MSG specification G is local-choice iff every choice node of G is local-choice.*

Local-choice MSG specifications have the communication after every choice node initiated by a single process — the choice leader. In [13] a deadlock-free realization with additional data in messages is proposed. It is easy to see that every MSG specification G is deadlock-free realizable if there is a local-choice MSG G' such that $\mathcal{L}(G) = \mathcal{L}(G')$. Note that the equivalence can be algorithmically checked due to [11].

Controllable specifications. The difficulties when realizing MSG are introduced by choice nodes. In local-choice MSG, the additional message content is used to ensure a single run through the graph is executed by all FSMs. In case of controllable-choice MSG, the additional content serves the same purpose but besides informing about the node the FSMs are currently executing the FSMs also attach a prediction about its future execution.

This allows us to relax the restriction on choice nodes and allows certain non-local choice nodes to be present in the specification. However, it is necessary to be able to resolve every occurrence of the choice node, i.e. make the decision in advance and inform all relevant processes.

Definition 7. Let $M = (E, <, \mathcal{P}, \mathcal{T}, P, \mathcal{M}, l)$ be a bMSC and $P' \subseteq \mathcal{P}$ be a subset of processes. A send event $e \in E$ is a resolving event for P' iff

$$\forall p \in P' \exists e_p \in P^{-1}(p) \text{ such that } e < e_p.$$

Intuitively, resolving events of M for P' can distribute information to all processes of P' while executing the rest of M , provided that other processes are forwarding the information.

Definition 8. Let $G = (S, \tau, s_0, s_f, L)$ be an MSG. A choice node u is said to be controllable-choice iff it satisfies both of the following conditions:

- For every path σ from s_0 to u there exists a resolving event in bMSC $L(\sigma)$ for $\text{triggers}(u)$.
- For every path $\sigma = s_1 s_2 \dots u$ such that $(u, s_1) \in \tau$, there exists a resolving event in bMSC $L(\sigma)$ for $\text{triggers}(u)$.

Intuitively, a choice node is controllable-choice, if every path from the initial node is labeled by a bMSC with a resolving event for all events initiating the communication after branching. Moreover, as it is necessary to attach only bounded information, the same restriction is required to hold for all cycles containing a controllable-choice node. In [8] we propose an algorithm that determines whether a given choice node is a controllable-choice node.

Definition 9. An MSG specification G is controllable-choice iff every choice node is either local-choice or controllable.

Note that there is no bound on the distance between the resolving event and the choice node it is resolving.

Local-choice vs. controllable-choice MSG. In the following, we show that the controllable-choice MSG are more expressive than local-choice MSG. It is easy to see that every local-choice MSG is also a controllable-choice MSG and that not every controllable-choice MSG is local-choice. In the following theorem, we strengthen the result by stating that the class of MSG that are language equivalent to some controllable-choice MSG is more expressive than the class of MSG that are language-equivalent to some local-choice MSG.

Theorem 1. The class of MSG that are language-equivalent to some local-choice MSG, forms a proper subset of MSG that are language-equivalent to some controllable-choice MSG.

Proof. Consider a MSG $G = (\mathcal{S}, \tau, s_0, s_f, L)$ with three nodes s_0, s_f and s , such that $(s_0, s), (s, s), (s, s_f) \in \tau$ and the only non-empty bMSC is $L(s)$ with two processes p, q . The projection of events on p is $p!q(m), p?q(m')$ and similarly for q the projection is $q!p(m'), q?p(m)$. Note that the only choice node s is controllable as both send events are resolving events for both of the processes.

The MSG G violates a necessary condition to be language equivalent to a local-choice specification. Intuitively, the condition states that its language must be a subset of a language of a generic local-choice equivalent MSG (for more details see [11]).

4 Realizability of Controllable-choice MSG

In this section we present an algorithm for realization of controllable-choice MSG. The class of local-choice specifications admits a natural deadlock-free realization because every branching is controlled by a single process.

As the *triggers* set for controllable-choice nodes can contain multiple processes, we need to ensure that all of them reach a consensus about which branch to choose. To achieve this goal, we allow the FSMs in certain situations to add a behavior prediction into its outgoing messages. Those predictions are stored in the finite-state control units and are forwarded within the existing communication to other FSMs.

The length of the prediction should be bounded, as we can attach only bounded information to the messages and we need to store it in the finite-state control unit. Therefore, it may be necessary to generate the behavior predictions multiple times. As the realization should be deadlock-free, we must ensure that the predictions are not conflicting — generated concurrently by different FSMs. To solve this we sometimes send together with the prediction also an event where the next prediction should be generated.

Definition 10. A prediction for an MSG $G = (\mathcal{S}, \tau, s_0, s_f, L)$ is a pair $(\sigma, e) \in \mathcal{S}^* \times (E \cup \perp)$, where E is the set of all events of bMSCs assigned by L , the path σ is called a prediction path, and e , called a control event, is an event from $L(\sigma)$. A prediction path must satisfy one of the following conditions:

- The prediction path σ is the longest common prefix of all MSG runs. This special initial prediction path is named *initialPath*.
- The prediction path σ is the shortest path $\sigma = \sigma_1\sigma_2 \dots \sigma_n$ in G satisfying
 1. $\sigma_n \in \mathcal{L}$, or
 2. $\sigma_n \in \mathcal{U} \wedge \exists 1 \leq i < n : \sigma_i = \sigma_n$, or
 3. $\sigma_n = s_f$,
 where $\mathcal{L} \subseteq \mathcal{S}$ is the set of all local-choice nodes and $\mathcal{U} \subseteq \mathcal{S}$ is the set of all controllable-choice nodes.

We refer to the first node and to the last node of a prediction path σ by *firstNode*(σ) and *lastNode*(σ), respectively.

Lemma 1. If the prediction path σ ends with a controllable-choice node u , the bMSC $L(\sigma)$ contains a resolving event for *triggers*(u) on $L(\sigma)$.

Proof. There are two cases to consider

- If $\sigma = \text{initialPath}$, then $\text{firstNode}(\sigma) = s_0$ and as node u is controllable-choice, the path σ contains a resolving event for $\text{triggers}(u)$.
- Otherwise, the controllable-choice node u occurs twice in the path σ . As every cycle containing a controllable-choice node has to contain a resolving event for the node, there is a resolving event for $\text{triggers}(u)$ on path σ .

As there are no outgoing edges allowed in s_f , the terminal node $s_f \notin \mathcal{U}$. \square

Note, that the number of events in a given MSG is finite and the length of each prediction path is bounded by $2 \cdot |\mathcal{S}|$.

When the CFM execution starts, every FSM is initialized with an initial prediction — $(\text{initialPath}, e_i)$ — and starts to execute the appropriate projection of $L(\text{initialPath})$. The value of e_i depends on the initialPath . Let $\text{lastNode}(\text{initialPath}) = \sigma_n$. In case of $\sigma_n \in \mathcal{U}$, the event e_i is an arbitrary resolving event from $L(\text{initialPath})$ for $\text{triggers}(\sigma_n)$. It follows from Lemma 1 that there exists such an event. If $\sigma_n \in \mathcal{L} \cup \{s_f\}$, we set $e_i = \perp$.

Every FSM stores two predictions, one that is being currently executed and a future prediction that is to be executed after the current one. Depending on the lastNode of the current prediction, there are the following possibilities where to generate the future prediction.

- If lastNode of the current prediction is in \mathcal{L} , the future prediction is generated by the local-choice leader, while executing the first event after branching.
- If lastNode of the current prediction is in \mathcal{U} , the future prediction is generated by an FSM that executes the control event of the current prediction, while executing the resolving event.
- If the lastNode of the current prediction is s_f , no further execution is possible and so no new prediction is generated.

When an FSM generates a new prediction, we require that there exists a transition in the MSG from the last node of the current prediction path to the first node of the future prediction path, as the concatenation of prediction paths should result in a path in the MSG. If an FSM generates a future prediction ending with a controllable-choice node u , it chooses an arbitrary resolving event for $\text{triggers}(u)$ to be the resolving event in the prediction. The existence of such an event follows from Lemma 1. To ensure that other FSMs are informed about the decisions, both predictions are attached to every outgoing message. The computation ends when no FSM is allowed to generate any future behavior.

4.1 Algorithm

In this section, we describe the realization algorithm. All the FSMs execute the same algorithm, an implementation of the FSM \mathcal{A}_p is described in Algorithm 1. We use an auxiliary function **path** that returns a prediction path for a given prediction. Every FSM stores a queue of events that it should execute — *eventQueue*. The queue is filled with projections of bMSCs labeling projection paths — $L(\text{prediction path})|_p$ for FSM \mathcal{A}_p . The execution starts with filling the queue with the projection of the *initialPath*.

Algorithm 1 Process p implementation

```

1: Variables:  $currentPrediction, nextPrediction, eventQueue$ ;
2:  $currentPrediction \leftarrow (initialPath, e_i)$ ;
3:  $nextPrediction \leftarrow \perp$ ;
4:  $eventQueue \leftarrow \text{push}(L(initialPath)|_p)$ ;
5: while true do
6:   if  $eventQueue$  is empty then
7:      $\text{getNextNode}()$ ;
8:    $e \leftarrow \text{pop}(eventQueue)$ ;
9:   if  $e$  is a send event then
10:    if  $e$  is the resolving event in  $currentPrediction$  then
11:       $node \leftarrow \text{lastNode}(\text{path}(currentPrediction))$ ;
12:       $nextPrediction \leftarrow \text{guessPrediction}(node)$ ;
13:       $\text{send}(e, currentPrediction, nextPrediction)$ ;
14:    if  $e$  is a receive event then
15:       $\text{receive}(e, cP, nP)$ ;
16:      if  $nextPrediction = \perp$  then
17:         $nextPrediction \leftarrow nP$ ;

```

Function 2 getNextNode function for process p

```

1: Function  $\text{getNextNode}()$ 
2:    $node \leftarrow \text{lastNode}(\text{path}(currentPrediction))$ ;
3:   if  $node \in \mathcal{U} \wedge p \in \text{triggers}(node)$  then
4:      $currentPrediction \leftarrow nextPrediction$ ;
5:      $nextPrediction \leftarrow \perp$ ;
6:      $eventQueue \leftarrow \text{push}(L(\text{path}(currentPrediction))|_p)$ ;
7:   else if  $node \in \mathcal{L} \wedge p \in \text{triggers}(node)$  then
8:      $currentPrediction \leftarrow \text{guessPrediction}(node)$ ;
9:      $nextPrediction \leftarrow \perp$ ;
10:     $eventQueue \leftarrow \text{push}(L(\text{path}(currentPrediction))|_p)$ ;
11:   else
12:      $currentPrediction \leftarrow \perp$ ;
13:      $nextPrediction \leftarrow \perp$ ;
14:      $\text{polling}()$ ;
15: end function

```

Function 3 Polling function for process p

```

1: Function  $\text{polling}()$ 
2:   while true do
3:     if  $p$  has a message in some of its input buffers then
4:        $\text{receive}(e, cP, nP)$ ;
5:        $currentPrediction \leftarrow cP$ ;
6:        $nextPrediction \leftarrow nP$ ;
7:        $eventQueue \leftarrow \text{push}(L(\text{path}(currentPrediction))|_p)$ ;
8:        $\text{pop}(eventQueue)$ ;
9:       return;
10: end function

```

The FSM executes a sequence of events according to its *eventQueue*. In order to exchange information with other FSMs, it adds its knowledge of predictions to every outgoing message, and improves its own predictions by receiving messages from other FSMs.

When the FSM executes a control event of the current prediction, it is responsible for generating the next prediction. The function **guessPrediction**(u) behaves as described in the previous section. It chooses a prediction (σ, e) , such that $(u, \text{firstNode}(\sigma)) \in \tau$. If $\text{lastNode}(\sigma) \in \mathcal{U}$, then e is a chosen resolving event in bMSC $L(\sigma)$ for the *triggers* set of the *lastNode*(σ). Otherwise, we leave $e = \perp$.

If the *eventQueue* is empty, the FSM runs the **getNextNode** function to determine the continuation of the execution. If the *lastNode* of the current prediction is a controllable-choice node and p is in the *triggers* set of this node, it uses the prediction from its variable *nextPrediction* as its *currentPrediction*. The variable *nextPrediction* is set to \perp .

If the *lastNode* of the *currentPrediction* is a local-choice node and p is the leader of the choice, it guesses the prediction and assigns it to the appropriate variables. Otherwise, the FSM forgets its predictions and enters a special polling state. This state is represented by the **Polling** function. Whenever the FSM receives a message, it sets its predictions according to the message. The pop function on line 8 ensures the consistency of the *eventQueue*.

An execution is finished successfully if all the FSMs are in the polling state and all the buffers are empty. The correctness proof of the following theorem is attached in the Appendix A.

Theorem 2. *Let G be a controllable-choice MSG. Then the CFM \mathcal{A} constructed by Algorithm 1 is a deadlock-free realization i.e. $\mathcal{L}(G) = \mathcal{L}(\mathcal{A})$.*

5 Conclusion

In this work we studied the message sequence graph realizability problem, i.e., the possibility to make an efficient and correct distributed implementation of the specified system. In general, the problem of determining whether a given specification is realizable is undecidable. Therefore, restricted classes of realizable specifications are in a great interest of software designers.

In recent years, a promising research direction is to study deadlock-free realizability allowing to attach bounded control data into existing messages. This concept turns out to be possible to realize reasonable specifications that are not realizable in the very original setting. In this work we introduced a new class of so called controllable-choice message sequence graphs that admits a deadlock-free realization with additional control data in messages. In other words, we have successfully extended the class of MSG conforming in the established setting of realizability. Moreover, we have presented an algorithm producing realization for a given controllable-choice message sequence graphs.

References

1. R. Alur, K. Etessami, and M. Yannakakis. Inference of message sequence charts. In *Proc. of ICSE*, pages 304–313, 2000.
2. R. Alur, K. Etessami, and M. Yannakakis. Realizability and verification of MSC graphs. *Theor. Comp. Science*, 331(1):97–114, 2005.
3. R. Alur, G.J. Holzmann, and D. Peled. An Analyzer for Message Sequence Charts. In *Proc. of TACAS*, LNCS, pages 35–48. Springer, 1996.
4. R. Alur and M. Yannakakis. Model Checking of Message Sequence Charts. In *Proc. of CONCUR*, number 1664 in LNCS, pages 114–129. Springer, 1999.
5. N. Baudru and R. Morin. Safe implementability of regular message sequence chart specifications. In *Proc. of ACIS*, pages 210–217, 2003.
6. H. Ben-Abdallah and S. Leue. Syntactic Detection of Process Divergence and Non-local Choice in Message Sequence Charts. In *Proc. of TACAS*, volume 1217 of LNCS, pages 259–274. Springer, 1997.
7. C.A. Chen, S. Kalvala, and J. Sinclair. Race Conditions in Message Sequence Charts. In *Proc. of APLAS*, volume 3780 of LNCS, pages 195–211. Springer, 2005.
8. M. Chmelík. Deciding Non-local Choice in High-level Message Sequence Charts *Bachelor thesis*, Faculty of Informatics, Masaryk University, Brno, 2009.
9. H. Dan, R.M. Hierons, and S. Counsell. A framework for pathologies of message sequence charts. *Information and Software Technology*, 2012. In press.
10. E. Elkind, B. Genest, and D. Peled. Detecting Races in Ensembles of Message Sequence Charts. In *Proc. of TACAS*, volume 4424 of LNCS, pages 420–434. Springer, 2007.
11. B. Genest and A. Muscholl. The Structure of Local-Choice in High-Level Message Sequence Charts (HMSC). Technical report, LIAFA, Université Paris VII, 2002.
12. B. Genest, A. Muscholl, and D. Kuske. A Kleene theorem for a class of communicating automata with effective algorithms. In *Proc. of DLT*, volume 3340 of LNCS. Springer, 2004.
13. B. Genest, A. Muscholl, H. Seidl, and M. Zeitoun. Infinite-State High-Level MSCs: Model-Checking and Realizability. *Journal of Computer and System Sciences*, 72(4):617–647, 2006.
14. A. Gill. *Introduction to the Theory of Finite-state Machines*. McGraw-Hill, 1962.
15. ITU Telecommunication Standardization Sector Study group 17. ITU recommendation Z.120, Message Sequence Charts (MSC), 2004.
16. C. Jard, R. Abdallah, and L. Hélouët. Realistic Implementation of Message Sequence Charts. Rapport de recherche RR-7597, INRIA, April 2011.
17. P.B. Ladkin and S. Leue. Interpreting Message Flow Graphs. *Formal Aspects of Computing*, 7(5):473–509, 1995.
18. M. Lohrey. Realizability of high-level message sequence charts: closing the gaps. *Theoretical Computer Science*, 309(1-3):529–554, 2003.
19. A.J. Mooij, N. Goga, and J. Romijn. Non-local choice and beyond: Intricacies of MSC choice nodes. In *Proc. of FASE*, volume 3442 of LNCS, pages 273–288. Springer, 2005.
20. A. Mousavi, B. Far, and A. Eberlein. The Problematic Property of Choice Nodes in high-level Message Sequence Charts. Tech. report, University of Calgary, 2006.
21. M. Mukund, K.N. Kumar, and M. Sohoni. Synthesizing distributed finite-state systems from MSCs. In *Proc. of CONCUR*, volume 1877 of LNCS, pages 521–535. Springer, 2000.
22. V. Řehák, P. Slovák, J. Strejček, and L. Hélouët. Decidable Race Condition and Open Coregions in HMSC. *Elect. Communications of the EASST*, 29:1–12, 2010.

A Correctness

Definition 11 ([1]). A word $w \in \Sigma^*$ is well-formed iff for every prefix v of w , every receive event in v has a matching send in v . A word $w \in \Sigma^*$ is complete iff every send event in w has a matching receive event in w .

Lemma 2. Let \mathcal{A} be a CFM and $w \in \mathcal{L}(\mathcal{A})$, then there exists a bMSC M such that $w \in \mathcal{L}(M)$.

Proof. Every $w \in \mathcal{L}(\mathcal{A})$ is a well-formed and complete word. Using results from [1] a word w is a bMSC (potentially non-FIFO) linearization iff it is well-formed and complete. So there exists a potentially non-FIFO bMSC M , such that $w \in \mathcal{L}(M)$. It remains to show, that the bMSC M satisfies the FIFO condition to fulfill our bMSC definition, but that follows directly from using FIFO buffers in the CFM. \square

Next, we make a few observations of the algorithm execution. For a given controllable-choice MSG G we construct a CFM $\mathcal{A} = (\mathcal{A}_p)_{p \in \mathcal{P}}$ according to Algorithm 1.

Lemma 3. Let (σ, e_i) be a prediction. FSM \mathcal{A}_p enters the **polling** function after executing $L(\sigma)$ iff

$$p \notin \text{triggers}(\text{lastNode}(\sigma)).$$

Proof. It holds for every prediction path σ that $\text{lastNode}(\sigma) \in \mathcal{U} \cup \mathcal{L} \cup \{s_f\}$. Note that $\text{triggers}(s_f) = \emptyset$ because no outgoing edge is allowed in the terminal state of an MSG. In case of $p \in \text{triggers}$, then $\text{lastNode}(\sigma) \in \mathcal{U} \cup \mathcal{L}$ and one of the two branches in Function 2 **getNextNode** is evaluated to true and **polling** function is skipped. \square

It is not necessarily true that every FSM executes an event in every prediction. In fact multiple predictions can be executed by the CFM, while a particular FSM \mathcal{A}_p executes the polling function and is not aware of predictions executed by other FSMs.

However, when a prediction path ends with a controllable-choice node, all the processes in the *triggers* set are active in the prediction.

Lemma 4. Let (σ, e_i) be a prediction, such that $\text{lastNode}(\sigma) \in \mathcal{U}$, then

$$p \in \text{triggers}(\text{lastNode}(\sigma)) \Rightarrow L(\sigma)|_p \neq \emptyset$$

Proof. Let $\text{lastNode}(\sigma) = u$. According to Lemma 1, there exists a resolving event for $\text{triggers}(u)$ in the bMSC $L(\sigma)$. Hence, there exists an event on process p that is dependent on the resolving event, therefore $L(\sigma)|_p \neq \emptyset$. \square

Another interesting observation is that it is possible to uniquely partition every MSG run into a sequence of prediction paths:

Proposition 2. Every run σ in G can be uniquely partitioned into a sequence of prediction paths such that $\sigma = \text{initialPath } w_2 \dots w_n$.

The following theorem shows that in fact it is not possible to execute simultaneously different predictions by different FSMs.

Theorem 3. *Let $\sigma = \text{initialPath } w_2 \dots w_n$ such that every w_i is a prediction path. Then every FSM \mathcal{A}_p for $p \in \text{triggers}(\text{lastNode}(w_n))$ possesses the same future prediction (w_{n+1}, e_{n+1}) , after executing the last event from $L(\sigma)$.*

Proof. We will prove the theorem by induction with respect to the length of path σ (measured by the number of prediction paths):

Base case Let the length of σ be 1, then $\sigma = \text{initialPath}$. We have to consider three options, depending on the type of the $\text{lastNode}(\text{initialPath})$:

- Let $\text{lastNode}(\text{initialPath}) = s_f$, then $\text{triggers}(\text{initialPath}) = \emptyset$ and there is nothing to prove.
- Let $\text{lastNode}(\text{initialPath}) \in \mathcal{L}$, then there exists a single leader process in the *triggers* set. The FSM representing the leader process may choose prediction (w_2, e_2) .
- The last option is that $\text{lastNode}(\text{initialPath}) \in \mathcal{U}$. Then the resolving event e_i in the initial prediction is not equal to \perp . The FSM executing the event guesses the next prediction (w_2, e_2) .

Let $p \in \text{triggers}(\text{lastNode}(\text{initialPath}))$. In case the FSM \mathcal{A}_p is not guessing the prediction, we need to show that it receives the prediction in some of its incoming messages. As e_i is a resolving event, there exists a dependent event on process p . Let us denote the minimal of such events e_p . Then e_p is a receive event and it is easy to see that the prediction (w_2, e_2) is attached to the incoming message. Hence, for every $p \in \text{triggers}(\text{lastNode}(\text{initialPath}))$, FSM \mathcal{A}_p has its variable *nextPrediction* set to (w_2, e_2) .

It follows from Lemma 3 that for every p not in the *triggers* set, the FSM \mathcal{A}_p is in the polling state having its variable *nextPrediction* set to \perp .

Induction step Let the length of σ be n . As in the base case, we have to consider multiple options:

- Let $\text{lastNode}(w_n) \in \{s_f\} \cup \mathcal{L}$, then the argument is the same as in the base case.
- So let $\text{lastNode}(w_n) \in \mathcal{U}$. From induction hypothesis, it follows that all FSMs \mathcal{A}_p for $p \in \text{triggers}(w_{n-1})$, start to execute prediction path w_n and all the others are in the polling state.

Let $p \in \text{triggers}(w_n)$. We show that FSM \mathcal{A}_p executes the projection $L(w_n)|_p$. It follows from Lemma 4 that this projection is non-empty. We have already shown that this is true for FSMs \mathcal{A}_p , such that $p \in \text{triggers}(w_{n-1})$. In case of $p \notin \text{triggers}(w_{n-1})$, the FSM \mathcal{A}_p is in the polling state. As it is not in the *triggers* set, its first action is a receive event. It is easy to see that the incoming message already contains the current prediction (w_n, e_n) and FSM \mathcal{A}_p starts to execute $L(w_n)|_p$.

The rest of the proof is similar to the base case. During the execution of the resolving event e_n a new prediction (w_{n+1}, e_{n+1}) is guessed and distributed to all FSMs \mathcal{A}_p for $p \in \text{triggers}(w_n)$.

□

To show that Algorithm 1 is a deadlock-free realization of the class of controllable-choice MSG we need to show that $\mathcal{L}(G) = \mathcal{L}(\mathcal{A})$. We will divide the proof into two parts, first showing that $\mathcal{L}(G) \subseteq \mathcal{L}(\mathcal{A})$ and finishing with $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(G)$.

A.1 $\mathcal{L}(G) \subseteq \mathcal{L}(\mathcal{A})$

We show that for all $w \in \mathcal{L}(G)$ also holds that $w \in \mathcal{L}(\mathcal{A})$. For every $w \in \mathcal{L}(G)$ there exists a run σ in G such that $w \in \mathcal{L}(L(\sigma))$.

We need to find a CFM execution, such that every FSM \mathcal{A}_p executes the projection $L(\sigma)|_p$ and ends in a polling state with the CFM having all the channels empty. Then using Proposition 1, follows $\mathcal{L}(M) \subseteq \mathcal{L}(\mathcal{A})$ and especially $w \in \mathcal{L}(\mathcal{A})$.

According to Proposition 2 we can partition every run σ uniquely into a sequence of prediction paths — *initialPath* $w_2 \dots w_n$. This sequence is a natural candidate for prediction paths that should be guessed during the CFM execution.

Every CFM execution starts with an initial prediction (*initialPath*, e_i). The guessed future prediction paths are w_2, w_3, \dots . The guessing continues until the last prediction path w_n is executed. As σ is a run in MSG G , $\text{lastNode}(w_n) = s_f$. Therefore, $\text{triggers}(\text{lastNode}(w_n)) = \emptyset$. It follows from Lemma 3 that all the FSMs are in the polling state. All the channels are empty because of the well-formedness and the completeness of the bMSC linearizations.

A.2 $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(G)$

We show that for every $w \in \mathcal{L}(\mathcal{A})$ also $w \in \mathcal{L}(G)$. According to Lemma 2, every $w \in \mathcal{L}(\mathcal{A})$ identifies a bMSC M . To conclude this part of the proof, we find a run σ in G , such that $M = L(\sigma)$. As $\mathcal{L}(M) \subseteq \mathcal{L}(G)$ we get $w \in \mathcal{L}(G)$.

The σ run in G is defined inductively. Every FSM starts with executing the *initialPath* prediction path. So it is safe to start the run σ with this prediction path.

According to Theorem 3 whenever some prediction w_i is executed, all FSMs \mathcal{A}_p for $p \in \text{triggers}(\text{lastNode}(w_i))$ agree on some future prediction w_{i+1} and all \mathcal{A}_p such that p executes an event in bMSC $L(w_{i+1})$, execute the projections $L(w_{i+1})|_p$. All the other FSMs are in the polling state and are awakened only if needed.

The predictions are guessed in such a way that the following condition holds:

$$(\text{lastNode}(w_i), \text{firstNode}(w_{i+1})) \in \tau$$

So it is safe to append w_{i+1} at the end of σ . Next we show that σ ends with a terminal node. The CFM accepts when all the channels are empty and all the FSMs are in the polling state. Hence, the last prediction that was executed ended with a node with an empty triggers set. In general it is possible that this is may not be the terminal node, but every path from this node reaches s_f without executing any event. So we can safely extend σ with a path to a terminal node.